AFRL-RI-RS-TR-2017-097

QUERY STORAGE AND RELAY IN RESEARCH ROOT (LACREND-RR)

UNIVERSITY OF SOUTHERN CALIFORNIA

MAY 2017

FINAL TECHNICAL REPORT

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED

STINFO COPY

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE

■ **AIR FORCE MATERIEL COMMAND** ■ **UNITED STATES AIR FORCE** ■ **ROME, NY 13441**

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09. This report is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (http://www.dtic.mil).

AFRL-RI-RS-TR-2017-097 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE CHIEF ENGINEER:

/ S /
FRANCES A. ROSE
Work Unit Manager

/ S /
JOHN D. MATYJAS
Technical Advisor, Computing
& Communications Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
MAY 2017	FINAL TECHNICAL REPORT	SEP 2015 – DEC 2016

4. TITLE AND SUBTITLE

Query Storage and Relay in Research Root (LACREND-RR)

5a. CONTRACT NUMBER
FA8750-15-2-0224

5b. GRANT NUMBER
N/A

5c. PROGRAM ELEMENT NUMBER
N/A

6. AUTHOR(S)

John Heidemann

5d. PROJECT NUMBER
DHSU

5e. TASK NUMBER
SA

5f. WORK UNIT NUMBER
RR

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
University of Southern California / Information Sciences Institute
4676 Admiralty Way, Ste. 1001
Marina del Rey, CA 90292

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Air Force Research Laboratory/RITE
525 Brooks Road
Rome NY 13441-4505

10. SPONSOR/MONITOR'S ACRONYM(S)
AFRL/RI

11. SPONSOR/MONITOR'S REPORT NUMBER
AFRL-RI-RS-TR-2017-097

12. DISTRIBUTION AVAILABILITY STATEMENT
Approved for Public Release; Distribution Unlimited. This report is the result of contracted fundamental research deemed exempt from public affairs security and policy review in accordance with SAF/AQR memorandum dated 10 Dec 08 and AFRL/CA policy clarification memorandum dated 16 Jan 09.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

This final report summarizes the objectives for the LACREND-RR project and the technical progress made against those objectives. The research objective of LACREND-RR is to support to capture and anonymize DNS data, and to provide a system that allows manipulation and replay of DNS data for experimental purposes.

15. SUBJECT TERMS

DNS, datasets, network security research

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			**FRANCES ROSE**
U	U	U	UU	21	19b. TELEPHONE NUMBER (*Include area code*) **N/A**

Contents

1 Summary

This final report summarizes the objectives for the LACREND-RR project and the technical progress made against those objectives. The research objective of LACREND-RR is to support to capture and anonymize DNS data, and to provide a system that allows manipulation and replay of DNS data for experimental purposes.

2 Introduction

The research objective of LACREND-RR is to support to capture and anonymize DNS data, and to provide a system that allows manipulation and replay of DNS data for experimental purposes.

3 Basic Programmatic Data

Performer: University of Southern California
Project Title: Query Storage and Relay in Research Root (LACREND-RR)
Agreement number: FA8750-15-2-0224
Period of Performance: 2015-09-04 to 2016-09-03
Estimated Total Award Value: $249k

3.1 Administrative Information

3.1.1 PI of record

John Heidemann, office telephone: +1 (310) 448-8708, e-mail address: johnh@isi.edu

3.1.2 Programmatic/Technical Reporter

Same

3.1.3 Administrative Contact

Jeanine Yamazaki, office telephone +1 (310) 448-8228, e-mail address: yamazaki@isi.edu

3.1.4 Financial Data Reporter

Joe Kemp, office telephone: +1 (310) 448-9171, e-mail address: kemp@isi.edu

3.1.5 Recipient Monitor

None

3.1.6 Subrecipient

None

3.2 Programmatic Information

3.2.1 Introduction: Project Description

3.2.1.1 Research Objectives

The research objective of LACREND-RR is to support to capture and anonymize DNS data, and to provide a system that allows manipulation and replay of DNS data for experimental purposes.

This effort contributes to focus area *FY16 Autonomous Defensive Cyber Operations* in BAA-AFRL-RIK-2015-0016 in several ways. Autonomous is a fundamental challenge Defensive Cyber Operations, with

the need to understand what actions can be automated and where human-in-the-loop is required. Ongoing research will advance that goal in these ways:

DNS as the defender: DNS systems today are often subject to Denial-of-Service attacks. Better methods are needed to harden and automate defensive actions DNS servers must take against such attacks (for example, DNS-over-TCP [Zhu15b]), to allow DNS services to reallocate resources nimbly in the face of an attack, and to understand how techniques such as anycast [Avramopoulos09a] interact with routing when under attack. Data from DNS and the ability to replay that data is essential to how autonomous defenses will react under attack and in regular operation.

DNS as part of the ecosystem: As part of the Internet, DNS is part of the attack/defense ecosystem. Attackers exploit DNS through amplification attacks [Kambourakis07a, Rossow14a], DNS injection attacks [herzberg13a], and as part of censorship [Duan12a, Anonymouas14a]. DNS data and experimental replay is needed to evaluate autonomous responses to these threats.

DNS as a sensor on attackers: DNS has proven a useful sensor on attackers for phishing and malware [Jung04a, Hao11a], and we have recently evaluate it for possible use in detecting a range of network wide events including scanners [Fukuda15a]. Sharable DNS data is needed to evaluate these techniques and how they may serve to predict attacks and trigger automated defenses.

Research in these methods requires the ability to collect, share, and experiment with DNS data. *This proposal will develop the capability to provide and experiment with DNS data.* This project will make data available through the LACREND project (http://www.isi.edu/ant/lacrend/) and the DHS PREDICT program (https://www.predict.org/). More importantly, this project will develop a new capability to replay DNS data against test software to advance the above research goals.

3.2.1.2 Public Problem Description

3.2.1.2.1 Public Research Goals/Contribution

The Domain Name System (DNS) is an important part of nearly every Internet transition, so correct operation of the DNS is essential. The DNS at risk from Denial-of-Service attacks, and subject to stress as it grows and evolves. Better tools are needed to support analysis of this system to improve its defense and evolution.

The research objective of LACREND-RR is to develop and support new approaches and tools that can observe and anonymize DNS data, and to provide a system that allows manipulation and replay of DNS data for experimental purposes.

3.2.1.2.2 Expected Impact

The primary expected impact of this work will be new tools to support analysis of the Domain Name System by academic, government, and commercial researchers of the Internet. Secondary impacts are that the work may generate suggestions to improve the DNS, or specific datasets about the DNS.

3.2.2 Technical Approach: Methods, Assumptions, and Procedures

3.2.2.1 Detailed Description of Public Technical Approach

The LACREND-RR project will develop new approaches to capture DNS traffic, anonymize it, and then replay it under experimental situations to evaluate new protocol and system designs. This work has three main components:

1. New approaches to capture DNS traffic and subject it to different kinds of processing.
2. New approaches to anonymize DNS traffic to allow its use and sharing while protecting query privacy.
3. New approaches to replay this traffic under controlled conditions.

3.2.2.2 Comparison with Current Technology

Current state-of-the-art relevant to this project is:

1. For DNS capture: general purpose capture systems like tcpdump (http://www.tcpdump.org) capture packets but do not explicitly support DNS. DNS-specific systems such as DSC (https://www.dns-oarc.net/oarc/data/dsc) capture DNS traffic but do not provide specific services for anonymization. We propose to build on the LANDER packet capture system (https://ant.isi.edu/software/lander/) and its ability to provide multiple queues of data streams, each with different users and levels of processing (raw packets, DNS streams, etc.) and anonymization (for clear to anonymized).
2. For DNS anonymization, we plan to build on techniques such as prefix-preserving IP anonymization. We will build on the existing dag_scrubber tool built at USC/ISI; we expect to build a new DNS-specific tool.
3. For DNS replay, we expect to build a new DNS replay system that can interoperate with live tools (DNS servers like bind and unbound). We will use testbeds such as DETER (http://deterlab.net) where it makes sense.

3.2.3 Schedule and Milestones

3.2.3.1 Schedule Graphic

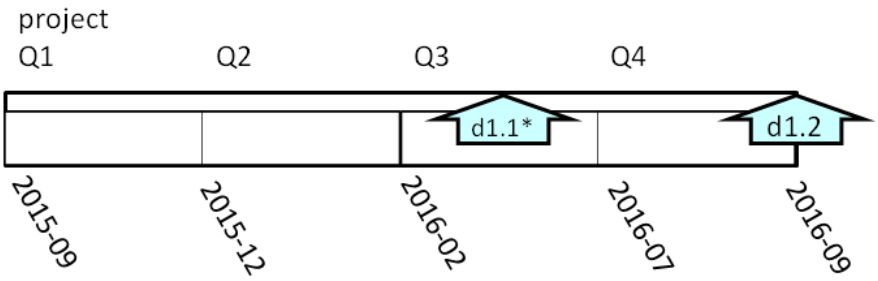

* Exact date of d1.1 subject to coordination with DNS-OARC; expected in Spring 2016.

3.2.3.2 Detailed Individual Task Descriptions

Task 1: Research Root Parallel Analysis and Playback. The contractor shall develop the following new capabilities to create datasets listed in Deliverables.

Subtask 1.1: The contractor shall develop a method of capture and archive of DNS queries that allows parallel analytic processing; allowing algorithms to run on commodity hardware at sustained rates of 100 Mb/s and bursts that go higher. *Completed by 2016-12-01 with [TT12][TT13]* .

Subtask 1.2: The contractor shall develop a method of DNS trace replay that allows stored data (using data from subtask 1.1, or potentially from other compatible datasets from PREDICT or elsewhere) to be fed into test software or hardware to evaluate performance. *Completed 2016-10-23 with [TT11]*

3.2.4 Deliverables Description

Deliverable 1.1: The contractor shall provide DNS data corresponding to the Day-in-the-Life-of-the-Internet experiments on an annual basis, to correspond with the measurement dates selected by DNS-OARC. *Completed on 2016-04-05*

Deliverable 1.2: The contractor shall provide the software developed as part of Subtask 1.2 as open source software, to be distributed on a public website. *Completed on 2016-10-24 with [TT11]*

Deliverables will be provided by USC/ISI by contract end-date.

3.2.5 Technology Transition and Technology Transfer Targets and Plans

Primary technology transition plan is (1) to work the B-Root operators and see our technology in use there. (2) To work with the LACREND project to provide tools and any developed datasets for distribution to the IMPACT program there.

In addition, where possible we expect to release our tools as open source software via the ANT website (https://ant.isi.edu/).

3.2.6 Data Rights

USC hereby grants to the U.S. Government a royalty free, world-wide, nonexclusive, irrevocable license to use, modify, reproduce, release, perform, display or disclose any data for Government purposes.

3.2.7 Quad Chart

Quad chart will be provided in a separate document.

4 Funding Report

Funding reports were provided monthly, separately from the quarterly and final technical reports.

5 Technical Report

5.1 Results and Discussion

5.1.1 Progress Against Planned Objectives

For details about progress towards planned objectives, please see all items listed in Section 4.1.2 marked [STx]. At this time we are working towards meeting all subtasks.

Task 1: Research Root Parallel Analysis and Playback.

.

Subtask 1.1: Method of capture and archive of DNS queries with parallel processing. *Completed by 2016-12-01 with [TT12][TT13]* .

Subtask 1.2: Method of DNS trace replay that allows stored data. *Completed 2016-10-23 with [TT11]*

5.1.2 Technical Accomplishments Over Contract

Accomplishments below are identified as pertaining to subtask [ST1.x], deliverable [D1.x], prototype improvement [PI] technology transfer [TTx], presentation [PreX], and publication [PubX].
This report includes discussion about some activities related to B-Root that are not directly funded by LACREND-RR. These activities are identified as [ext] and are provided to give context for the LACREND-RR work. They are supported by other (non-LACREND-RR) funding.

1. **[ext; PI] B-Root Infrastructure/Architecture:** On 2015-08-07, just before the contract began we deployed a completely new architecture for B-Root. This infrastructure was rebuilt from the ground up by our operations team to use centrally managed infrastructure (controlled by Puppet) and new hardware. The immediate goals were (a) to improve IPv6 performance by providing more than 1 IPv6 server, (b) to update the hardware and software to modern systems and versions, (c) to centralize management to make it maintainable by a small group, and (d) to improve overall performance. All of those objectives were met. On 2015-08-07 IPv6 went live on the new architecture, and on 2015-08-12 IPv4 went live as well.
2. **[ext; PI] B-Root Infrastructure/Architecture:** On 2015-12-16 we deployed a firewall in front of the B-LAX site and upgraded incoming network bandwidth to 10 Gb/s. These changes are a response to the events of 2016-11-30, where many roots received 100x normal traffic and B-Root suffered serious query loss. We believe our 10 Gb/s firewall would have allowed us to tolerate this DoS event.
3. **[ext; PI] B-Root Infrastructure/Architecture:** In 2016q1 we defined our next generation (2016) architecture for B-Root: a pair of front-end nodes dual linked to four back-end nodes. One front-end will be dedicated as a firewall and router, the other as measurement. All network connections will be 10 Gb/s. Our goal is to deploy this architecture at our second site. In addition, we believe we can migrate our current hardware at the LAX site to provide a mini-version of this architecture, although with only 1 Gb/s internal links. In 2016q1 Michael Elkins began prototyping this 2016 architecture in our lab, demonstrating functioning 10Gb/s and software routing.
4. **[ext; PI] B-Root Infrastructure/Architecture:** In 2016q2 Michael Elkins continued prototyping our 2016 architecture in our lab, demonstrating an application-specific load balancing and working out 10 Gb/s OCSP for back-end load balancing. Surprisingly, IPv6 OCSP requires a very recent Linux kernel (4.4).
5. **[ext; PI] B-Root Infrastructure/Architecture:** In 2016q3 Michael Elkins continued prototyping our 2016 architecture in our lab, finding an ECMP implementation that we can use with current Linux.
6. **[ext; PI] B-Root Infrastructure/Architecture:** In 2016q4 we validated our 2016 architecture with the upgrade of LAX to a proto-2016 state, as described in the next item.

7. **[ext; PI] B-Root Infrastructure/Configuration:** In 2016-07 we evaluated that our configuration management software in Puppet needed major revision. We decided to switch tools as part of this process to ansible, a more modern tool. We expect this transition to take the next 9 months.
8. **[ext; PI] B-Root Infrastructure/Configuration:** In 2016q4 we continued our migration of B-Root from Puppet to Ansible. Work is nearing completion, and we expect to finalize this transition by 2017q1.
9. **[ext; PI] B-Root Infrastructure/Operations:** On 2016-02-23 we lost one backend node due to hardware failure. With multiple backends, external performance was not affected and we operated with one fewer back-end for about 4 weeks until we replaced the failed box on 2016-03-19.
10. **[ext; PI] B-Root Infrastructure/Operations:** In 2016q1 Wes Hardaker deployed Nagios alterting. Nagios will inform us (via e-mail and SMS) when there are serious changes in the B-Root operational status.
11. **[ext] B-Root Infrastructure/Operations:** In 2016q2, following several DoS attacks of different sizes, we began cataloging responses to attacks as a "playbook" with canned methods for detection, filtering, etc. We expect to evolve this playbook over coming years.
12. **[ext] B-Root Infrastructure/Operations:** In 2016-06-25 there was a large DDoS attack on all Root Letters. B-Root performed better than in the 2015-11 attack, but we did not get the benefit of our full inbound bandwidth because of delays processing packets resulted in backpressure on the upstream router, and because of congestion in our upstream ISP. We fixed the backpressure issue so we expect to perform better in the next attack.
13. **[ext] B-Root Infrastructure/LAX Site:** Because of continued delays deploying our second site, in 2016q3 we committed to upgrading the LAX site to support greater outbound capacity. We expect this upgrade to happen by 2016-10-30.
14. **[ext] B-Root Infrastructure/LAX Site:** Because of continued delays deploying our second site, in 2016q4 we committed to upgrading the LAX site to support greater outbound capacity. On 2016-11-23, Michael Elkins completed transition of our LAX site to a simpler version of our 2016 architecture—one computer serves and load balancer and firewall, and it is directly connected to 4 backends, and a final computer does data collection. We expect to continue to refine this architecture to support firewall/measurement failover in 2017, but we are very happy to have transitioned away from dedicated routers
15. **[ext] B-Root Infrastructure/Second Site:** The B-Root operations team has been discussing deployment to a second site. We expect to deploy sometime in 2016 or 2017, and we will bring the LACREND-RR measurement infrastructure forward to encompass a second site. In 2016q1 we discussed potential sites and, after weighing alternatives, identified Miami as the best location fitting our constraints. We expect to be hosted by FIU/Amlight.net based on early discussions. A formal contract awaits equipment purchase.
16. **[ext] B-Root Infrastructure/Second Site:** The B-Root operations team has been discussing deployment to a second site. In 2016q2, deployment plans for the second site are on hold pending resolution of USC funding. We had several internal meetings at USC to make our case and will not know status until 2016-07-01.
17. **[ext] B-Root Infrastructure/Second Site:** The B-Root operations team has been discussing deployment to a second site. In 2016q3 we committed to a location (Miami) and began negoations on a hosting agreement, and plans to purchase hardware. We expect to carry this deployment out by the end of the 2016.
18. **[ext] B-Root Infrastructure/Second Site:** The B-Root operations team has been discussing deployment to a second site. In 2016q4 we completed contracting agreements with Amlight for our second site in Miami. In 2015q4 the B-Root team discussed possible locations for a second site. We expect to deploy in 2017q1, as soon as we complete software installation.
19. **[ext] B-Root Infrastructure/Address Space:** In 2015-09 John Heidemann worked with ARIN to allocate an AS number and a /23 IPv4 address block for B-Root use. We were fortunate to get

some of the last ARIN IPv4 address space. We also identified TI has holding the /24 that is adjacent to our current operational /24 (so we can form a /23). We hope to swap for that address space.

20. **[ext; PI] B-Root Infrastructure/Address Space:** In preparation for a 2nd site, in 2016q1, Suzanne Woolf started discussions with ARIN and TI about swapping the TI IPv4 /24 prefix that is adjacent to our current operational /24 so we can form a /23.
21. **[ext; PI] B-Root Infrastructure/Address Space:** In preparation for a 2nd site, in 2016q2 Suzanne Woolf continues discussions with ARIN and TI about swapping the TI IPv4 /24 prefix that is adjacent to our current operational /24 so we can form a /23. In 2016q2 we determined that the best path to proceed is to directly execute an agreement between USC and TI, then follow up with ARIN with this documentation.
22. **[ext; PI] B-Root Infrastructure/Address Space:** In preparation for a 2nd site, in 2016q3 Suzanne Woolf continues discussions with ARIN and TI about swapping the TI IPv4 /24 prefix that is adjacent to our current operational /24 so we can form a /23. In 2016q3 we completed USC's side of the document swapping addresses with TI, but we are waiting on TI for their side.
23. **[ext; PI] B-Root Infrastructure/Address Space:** In preparation for a 2nd site, in 2016q4 Suzanne Woolf continues discussions with ARIN and TI about swapping the TI IPv4 /24 prefix that is adjacent to our current operational /24 so we can form a /23. In 2016q4 we moved forward on the TI side of address space transfer. While previously they were willing to exchange the address space for goodwill, they now with to review considerations. We expect this discussion will continue in 2017 and we will consider IPv4 renumbering as a backup blan.
24. **[ext; PI] B-Root Infrastructure/Address Space:** In preparation for a 2nd site, we decided to renumber IPv6, since adjacent IPv6 space is occupied by others and unavailable. In 2016q3, John Heidemann acquired new IPv6 address space for use and we plan to deploy it this fall.
25. **[ext; PI] B-Root Infrastructure/Address Space:** In preparation for a 2nd site, we decided to renumber IPv6, since adjacent IPv6 space is occupied by others and unavailable. In 2016q4 we began routing plans to deploy this new address space, and we expect to activate it in 2017q1.
26. **[ext] B-Root Infrastructure/Routing:** In 2015-10 we began plans to have B-root begin routing from its own AS. This step must be completed before we have a second anycast site. We expect this task to take several months.
27. **[ext; PI] B-Root Infrastructure/Routing:** In 2014q1, John Heidemann worked with ARIN to allocate new AS394353 for eventual use with routing.
28. **[ext] B-Root Infrastructure/Testbed:** We maintain a tested (a mini-version of B Root) to experiment with new hardware and software configurations. In 2015-08 we shipped the testbed into deployment. In 2015-10 Michael Elkins rebuilt the testbed so we can continue development.
29. **[ext] B-Root Infrastructure/Testbed:** We maintain a testbed (a mini-version of B Root) to experiment with new hardware and software configurations. In 2016q1 we began to redeploy our testbed, taking steps to isolate it from the Internet with a firewall so it can run with operational addresses without interfering with actual operation.
30. **[ext] B-Root Infrastructure/Testbed:** In 2016q2 we continued incrementally building out our testbed to match our 2016 architecture. Michael Elkins added 10Gb/s internal links and an optical splitter, and Yuri Pradkin deployed LANDER packet capture using the 10 Gb/s optical splitter and added cacti monitoring to allow testbed use and observation.
31. **[ext] B-Root/Outreach:** In 2015-10 we discussed providing background about the history of B-Root as a courtesy to RSSAC.
32. **[TT1][ext] B-Root/Outreach:** We discussed providing background about the history of B-Root as a courtesy to RSSAC. In 2016q1 we iterated on this document with RSSAC; we expect it to go out for review in the RSSAC caucus in 2016q2. B-Root is also working with RSSAC on several other issues, currently for internal RSSAC discussion.

33. **[TT1][ext] B-Root/Outreach:** In 2016q4 Wes Hardaker and Suzanne Woolf have continued to collaborate with RSSAC about topics related to Root Operations, with meetings in Hyderabad, India at ICANN, and Seoul, Korea at IETF.
34. **[TT5][Pub9][ext] B-Root/Outreach-DINR:** As part of the research root we are planning a workshop, DNS and Internet Naming Research Directions 2016 (https://ant.isi.edu/events/dinr2016/) , was held at USC/ISI on 2016-11-17. With 61 authors and attendees and 21 talks, we are very happy with the outcome and interest in DNS research.
35. **[TT5][ext] B-Root/Outreach-DINR:** As part of the research root we are planning a workshop, DNS and Internet Naming Research Directions 2016 (https://ant.isi.edu/events/dinr2016/) , to be held at USC/ISI on 2016-11-17. In 2016q2, John Heidemann added David Dagon as co-PI of the workshop.
36. **[TT5][ext] B-Root/Outreach-DINR:** As part of the research root we are planning a workshop, DNS and Internet Naming Research Directions 2016 (https://ant.isi.edu/events/dinr2016/) , to be held at USC/ISI on 2016-11-17. In 2016q3, John Heidemann and David Dagon worked on PR for the workshop and added Mark Allman as an additional co-chair.
37. **[TT5][ext] B-Root/Outreach-DINR:** As part of the research root we are planning a workshop, DNS and Internet Naming Research Directions 2016 (https://ant.isi.edu/events/dinr2016/) , to be held at USC/ISI on 2016-11-17. John Heidemann announced this workshop at CAIDA AIMS on 2016-02-10..
38. **[ST1.1; PI] Data Collection/Infrastructure:** As part of our 2015 architecture, in 2015-10 we have deployed a new data collection system. This system is based on the LANDER packet collection system (currently maintained by the LACREND project) and is currently including full packet captures and archive with the LACREND project. We plan to expand this system in coming months to facilitate data sharing. The LANDER deployment was the responsibility of Yuri Pradkin and Abdul Qadeer.
39. **[ST1.1; PI] Data Collection/Infrastructure:** Data collection for B-Root consists of the LANDER packet capture software, a hadoop cluster for local processing, and data archive at ISI. In 2016q1 we restored one failed Hadoop node and activated one additional node in this cluster. Also in 2016q1, Abdul Qadeer refined and tuned the data processing software pipeline.
40. **[TT6] Data Collection/anonymization:** In 2016-01 we released *dnsanon*, a new program that parses pcap traces and outputs dns data, with optional anonymization. We expect to expand on this as part of our RSSAC-002 processing stack.
41. **[TT6][TT7][PI] Data Collection/RSSAC:** In 2016q2, John Heidemann joined the RSSAC work party planning to revise the RSSAC-002 statistics specification. Concurrent with this he implemented a new processing stack to do RSSAC-002 statistics. Instead of a centralized database that requires significant resources, our stack uses Hadoop, with a new pcap parser (dnsanon) and RSSAC analyzer (dnsanon_rssac). We released updates to dnsanon several times in 2016q2, and did a public release of dnsanon_rssac in 2016-05 at https://ant.isi.edu/software/dnsanon_rssac.
42. **[TT6][TT7][TT9][PI] Data Collection/RSSAC:** In 2016q2, John Heidemann joined the RSSAC work party planning to revise the RSSAC-002 statistics specification. Concurrent with this he implemented a new processing stack to do RSSAC-002 statistics. Instead of a centralized database that requires significant resources, our stack uses Hadoop, with a new pcap parser (dnsanon) and RSSAC analyzer (dnsanon_rssac). On 2016-05-28 we deployed this tool for production use at B-Root.
43. **[TT6][TT7][TT9][TT14][PI] Data Collection/RSSAC:** In 2016q4, John Heidemann revised his RSSAC-002 processing software to fix some bugs in multi-site statistics, with a new release of dnsanon_rssac-1.4 on 2016-12-11 and dnsanon_rssac-1.5 on 2016-12-29. These changes will support multi-site B-Root.
44. **[TT4] Data Collection/RSSAC:** In 2016-12 Abdul Qadeer deployed Hedgehog, the ICANN-produced processing stack to produce RSSAC-002 statistics. Deployment was somewhat painful,

with a large back-end database that required tuning. In addition, we found an incompatibility between the Hedgehog design (or a widely ancasted service with regular pcap files) and our deployment (a single site with frequent pcap files not aligned to minute boundaries) that results in slightly to statistics. In the long-run we hope to replace Hedgehog with our own RSSAC-002 processing stack.

45. **[D1.1] Data Collection/DITL:** In 2016q1, DNS-OARC has identified 2016-04-05 as the target date for DITL this year. We expect DITL to occur as part of our regular data collection activities, and Abdul Qadeer, Yuri Pradkin, and John Heidemann reviewed and refreshed our DITL handling procedure.
46. **[D1.1] Data Collection/DITL:** On 2016-04-05 we collected data as part of the 2016 DITL event, working with DNS-OARC and the other root operators.
47. **[Pre20] [Pre22][Pre23] Data collection/Testbed:** Wes Hardaker gave a talk about our plans for a DNS testbed at the ICANN DNSSEC Workshop in Hyderabad, India on 2016-11-07. John Heidemann gave a similar talk at DINR-2016 on 2016-11-17. John Heidemann talked about this in part of his talk on DDoS datasets at ACSAC 2016 on 2016-12-07.
48. **[ST1.1; PI] Data Collection/Processing**: To support analysis of B-Root data we have deployed a new 16-node Hadoop cluster at the same location as the 2015 B-Root operational site. Currently this infrastructure is in operational use to archive packet data. This deployment was by Yuri Pradkin with data movement and processing by Abdul Qadeer.
49. **[ST1.1; PI] Data Collection/Processing**: The software making up B-Root data processing includes file conversion from pcap to text format, and compression and archive in both formats. In 2016q1, Abdul Qadeer revised this software pipeline. In 2016q1 John Heidemann deployed software to allow data sharing with other parties, where a window of data is made available from the archive.
50. **[ST1.1; PI] Data Collection/Processing**: The software making up B-Root data processing includes file conversion from pcap to text format, and compression and archive in both formats. In 2016q2, John Heidemann and Abdul Qadeer revised this software pipeline to cope with week-long 2x load following the AirOS worm, starting 2016-04-21. We were able to improve throughput in our Hadoop cluster to handle the load.
51. **[TT12][TT13] [ST1.1] Data Collection/Processing:** In 2016q4, Abdul Qadeer has implemented our current DNS processing architecture using a combination of LANDER packet capture and Hadoop parallel processing. Based on what we have learned, Abdul Qadeer and John Heidemann have developed a design for a new data processing architecture using Apache Tez that should greatly improve parallelism and lower latency. We expect to implement this design in future work.
52. **[ST1.2; PI] Data Collection/Trace Replay:** Liang Zhu has begun development of a DNS trace replay system. In 2015q4 he has developed replay software with a multi-level DNS system to recreate the recursive resolver infrastructure in the lab. The next piece is to add software to generate representative zone files; we expect that work to be prototyped in 2016q1.
53. **[ST1.2; PI] Data Collection/Trace Replay:** Liang Zhu has continues development of a DNS trace replay system. In 2016q1 he completed software to generate representative zone files from traces. He expects to begin evaluation of his trace replay system in 2016q2.
54. **[ST1.2; PI] Data Collection/Trace Replay:** Liang Zhu has begun development of a DNS trace replay system. In 2016q2 he transitioned a version of this software for use in our testbed, where it was used by Michael Elkins and Yuri Pradkin for load testing. Liang Zhu continues to evaluate his software to measure its accuracy.
55. **[TT11][TT11] [ST1.2; PI] Data Collection/Trace Replay:** Liang Zhu has begun development of a DNS trace replay system. In 2016q3 he continued to evaluate this tool for accuracy, and we began work on a paper on it. Initial release of his replay system, LDplayer/dns-replay-controller-0.1, occurred on 2016-10-24.

56. **[TT11][TT11] [ST1.2; PI] Data Collection/Trace Replay:** In 2016-10, Liang Zhu continues development of a DNS trace replay system. Initial release of his replay system, LDplayer/dns-replay-controller-0.1, occurred on 2016-10-24, and he continues to refine it.
57. **[D1.1] Data Collection/Events:** In 2015-09 we began an internal log to document specific data events pertaining to B-Root, such as large Denial-of-Service attacks. These events may turn in to curated datasets in the future.
58. **[TT1] Data Collection/Events:** In 2015-10 we revised our operational practices to include Receiver Rate Limiting, based on data collected as part of LACREND-RR.
59. **[TT2]**Error! Reference source not found. **Data Collection/Events:** On 2015-11-30 multiple root operators suffered a very large DDoS attack. B-Root suffered high loss rates during most of the attack. This attack was significant enough that the Root Operators documented it in a memo (Events of 2015-11-30, at http://www.root-servers.org/news/events-of-20151130.txt). Our response to the attack was to plan deployment of a firewall and a 10 Gb/s connection to our ISP, to be deployed in 2015-12. We believe this deployment would have stopped this specific attack**.**
60. **[D1.1] Data Collection/Events:** We collect notes about events that affect B-Root. On 2016-12-27 we observed a 2.5 Gb/s traffic surge that was blocked with our recently deployed firewall.
61. **[ext][TT3][Pub1][Pre5] Research/DNS Backscatter:** On 2015-10-29 Kensuke Fukuda presented the paper "Detecting Malicious Activity with DNS Backscatter" at ACM IMC 2016 in Tokyo, Japan. Work on this paper pre-dates LACREND-RR, but the paper made use of a prototype of the facilities we are developing in this effort.
62. **[TT3] [Pub1][Pre5] Research/DNS Backscatter:** In 2016q1, Abdul Qadeer worked with John Heidemann and Kensuke Fukuda to revise the work in DNS backscatter to better understand how the machine learning approaches described in the IMC 2016 paper behave over time (as the underlying data changes). We hope to use this addition as part of a journal submission in 2016q2.
63. **[TT8][Pub5][Pre12][ext] Data Collection/Anycast:** on 2016-05-18, we released the technical report "Anycast Latency: How Many Sites Are Enough?" by Ricardo de O. Schmidt, John Heidemann and Jan Harm Kuipers. This report summarizes our the impact of physical location on anycast and has influenced the location of B-Root's second site. This paper is also currently under submission for a peer-reviewed conference. In 2016-06 we found that the paper was rejected, so we plan to revise and resubmit to another conference this fall.
64. **[TT8][Pub5][Pre12][Pub8][ext] Research/Anycast:** In 2016q4 we revised our paper "Anycast Latency: How Many Sites Are Enough?" by Ricardo de O. Schmidt, John Heidemann and Jan Harm Kuipers and resubmitted it to PAM (the Passive and Active Measurements Conference). The paper was accepted to appear in Sydney, Australia in March 2017.
65. **[TT8][Pub5][ext] Research /Anycast:** on 2016-05-18, we released the technical report "Anycast Latency: How Many Sites Are Enough?" by Ricardo de O. Schmidt, John Heidemann and Jan Harm Kuipers. This report summarizes our the impact of physical location on anycast and has influenced the location of B-Root's second site. This paper is also currently under submission for a peer-reviewed conference.
66. **[Pre17][Pre18] Research /Anycast:** In addition, several people gave several talks summarizing the paper "Anycast Latency: How Many Sites Are Enough?" by Ricardo de O. Schmidt, John Heidemann and Jan Harm Kuipers at several events: John Heidemann at the ICANN Lunch Seminar on 2016-10-11; John Heidemann at DNS-OARC on 2016-10-16.
67. **[TT2][Pub6] [ext] Data Collection/DDoS:** On 2016-05-18, we released the technical report "Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event (extended)" by Giovane C. M. Moura, Ricardo de O. Schmidt, John Heidemann, Wouter B. de Vries, Moritz Müller, Lan Wei and Christian Hesselman. This paper is also currently under submission for a peer-reviewed conference.
68. **[TT2][Pub6] [Pre13][Pre14][Pre15][ext] Data Collection/DDoS:** On 2016-05-18, we released the technical report "Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event (extended)" by Giovane C. M. Moura, Ricardo de O. Schmidt, John Heidemann, Wouter B. de

Vries, Moritz Müller, Lan Wei and Christian Hesselman. This paper was accepted for appearance in ACM IMC in 2016-11; in 2016q3 we are revising it..

69. **[TT2][Pub6] [Pre13][Pre14][Pre15][Pub7][ext] Research /DDoS:** On 2016-11-15, Giovane Moura presented our paper "Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event (extended)" by Giovane C. M. Moura, Ricardo de O. Schmidt, John Heidemann, Wouter B. de Vries, Moritz Müller, Lan Wei and Christian Hesselman at the ACM Internet Measurements Conference in Santa Monica, California.
70. **[Pre16][Pre19] Research/Anycast:** In addition, in 2016q4 several people gave several talks summarizing the paper "Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event (extended)" by Giovane C. M. Moura, Ricardo de O. Schmidt, John Heidemann at several events: John Heidemann at the USC NSL lunch meeting on 2016-09-13; John Heidemann at DNS-OARC on 2016-10-16.
71. **[ext] B-Root/Personnel:** In 2015-11 Jim Kodak, long-term B-Root operator, stepped down. He was not involved in the new architecture and his transfer to other duties at ISI was long planned but we will miss his 20 years of experience with B-Root.

5.1.3 Improvements to Prototypes Over Contract

For specific improvements to prototypes in this period, please see all items listed in Section 4.1.2 marked [PI].

5.1.4 Significant Changes to Technical Approach Over Contract

Over the contract, we had no significant changes to technical approach from what was proposed.

5.1.5 Deliverables Over Contract

For progress against deliverables in this period, please see all items listed in Section 4.1.2 marked [Dx]. At this time we are working towards meeting all deliverables.

Deliverable 1.1: DNS data for Day-in-the-Life-of-the-Internet. *Competed in 2016-04*

Deliverable 1.2: Release software developed as part of Subtask 1.2. *Completed 2016-10-23 with [TT11]*

5.1.6 Technology Transition and Transfer this Period

Technology transfer is to include description, list, and contacts. Unless otherwise indicated, all technology transfer is to B-Root Operations, with Terry Benzel (USC/ISI) as the point-of-contact.

For technology transition in this period, please see all items listed in Section 4.1.2 marked [TTx].

Summary of new technology transfer events:

[TT1] Evaluation of BIND DNS RRL (Receiver Rate Limiting) and sharing this information with the community.
[TT2] Description of the 2016-11-30 event that affected DNS Root Server performance.
[TT3] Use of B-Root data to evaluate DNS backscatter

New in 2016q1:

[TT4] B-Root history provided to RSSAC history report
[TT5] Planned DNS and Internet Naming Research Directions 2016 (https://ant.isi.edu/events/dinr2016/) workshop.
[TT6] Public release of dnsanon program (pcap-to-text and anonymizer, at https://ant.isi.edu/software/dsnaon/)

New in 2016q2:

[TT7] Public release of dnsanon_rssac, taking dnsanon output to produce RSSAC-002 statistics.

[TT8] Evaluation of the effects of anycast location on latency and sharing this information with the community.

New in 2016q3:

[TT9] Updates to dnsanon and dnsanon_rssac based on use.
[TT10] Release of LDplayer/replayer-0.1 at https://ant.isi.edu/software/ldplayer/
[TT11] Release of dns trace replay software prototype.

New in 2016q4:

[TT12] Developed parallel processing architecture for DNS processing using LANDER and Apache Hadoop.
[TT13] Developed design of improved processing architecture for DNS processing using Apache Tez (building on our Hadoop infrastructure).
[TT14] Updates to dnsanon_rssac to handle multiple sites. https://ant.isi.edu/software/dnsanon_rssac/

5.1.7 Publications Over Contract

Publications are to include Title, Author, date, venue, and keywords. For context about these publications, please see items listed in Section 4.1.2 marked [PubX].

[Pub1] Kensuke Fukuda and John Heidemann. Detecting Malicious Activity with DNS Backscatter. In Proceedings of the ACM Internet Measurement Conference, pp. 197-210. Tokyo, Japan, ACM. October, 2015. <http://dx.doi.org/10.1145/2815675.2815706>, <http://www.isi.edu/%7ejohnh/PAPERS/Fukuda15a.html>.
[Pub2] Rootops. Events of 2015-11-30. Technical Report Root Server Operators, Dec. 4, 2015. <http://www.root-servers.org/news/events-of-20151130.txt>.

New in 2016q1:

[Pub3] John Heidemann. New Opportunities for Research and Experiments in Internet Naming And Identification. Talk at Active Internet Measurement Workshop. http://www.isi.edu/~johnh/PAPERS/Heidemann16a.html

New in 2016q2:

[Pub4] Z. Hu, L. Zhu, J. Heidemann, A. Mankin, D. Wessels and P. Hoffman. Specification for DNS over Transport Layer Security (TLS) . RFC-7858. Internet Request For Comments. http://dx.doi.org/10.17487/RFC7858
[Pub5] Ricardo de O. Schmidt, John Heidemann and Jan Harm Kuipers. Anycast Latency: How Many Sites Are Enough? Technical Report ISI-TR-2016-708. USC/Information Sciences Institute. http://www.isi.edu/~johnh/PAPERS/Schmidt16a.pdf
[Pub6] Giovane C. M. Moura, Ricardo de O. Schmidt, John Heidemann, Wouter B. de Vries, Moritz Müller, Lan Wei and Christian Hesselman. Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event (extended). Technical Report ISI-TR-2016-709. USC/Information Sciences Institute. http://www.isi.edu/~johnh/PAPERS/Moura16a.pdf

New in 2016q3:

No new publications, but the Nov. 30 paper was accepted to appear at ACM IMC 2016.

New in 2016q4:

[Pub7] Giovane C. M. Moura, Ricardo de O. Schmidt, John Heidemann, Wouter B. de Vries, Moritz Müller, Lan Wei and Christian Hesselman. Anycast vs. DDoS: Evaluating the November 2015 Root DNS Event. In *Proceedings of the ACM Internet Measurement Conference*, November, 2016. <http://dx.doi.org/10.1145/2987443.2987446>, <http://www.isi.edu/%7ejohnh/PAPERS/Moura16b.html>.
[Pub8] Ricardo de O. Schmidt, John Heidemann, and Jan Harm Kuipers. Anycast Latency: How Many Sites Are Enough?. In _Proceedings of the Passive and Active Measurement Workshop_, p. to appear. Sydney, Australia, Springer. May, 2017. <http://www.isi.edu/%7ejohnh/PAPERS/Schmidt17a.html>.
[Pub9] DNS and Internet Naming Research Workshop (DINR-2016) website. https://ant.isi.edu/events/dinr2016/

5.1.8 Meetings and Presentations Over Contract

Meetings and presentations are to include meeting name, purpose, dates, location, attendees, and name of the presentation.
For context about these presentations listed here, please see items listed in Section 4.1.2 marked [PreX].

[Pre1] On 2015-08-27 John Heidemann and Terry Benzel met with David Galassi, deputy CIO at USC in charge of infrastructure, to update him about the progress in B-Root and to discuss USC support for B-Root.
[Pre2] We are planning for 2015-11-01, where John Heidemann and Wes Hardaker will make a presentation about the new B-Root architecture to Root-Ops.
[Pre3] On 2015-09-23 to -24, Wes Hardaker and Suzanne Woolf attended the RSSAC Retreat on behalf of B-Root. RSSAC retreats discuss Root Letter responsibilities and expectations.
[Pre4] On 2015-10-15, Suzanne Woolf attended an ICANN meeting in Dublin, Ireland, partially on behalf of B-Root.
[Pre5] On 2015-10-29 Kenuske Fukuda presented the paper “Detecting Malicious Activity with DNS Backscatter” at ACM IMC 2016.

[Pre6] On 2015-11-20, Terry Benzel and John Heidemann met with the USC acting CIO David Shook to update him about the progress in B-Root and to discuss USC support for B-Root.

New in 2016q1:

[Pre7] On 2015-02-02 John Heidemann presented gave the talk "Anycast Latency: How Many Sites are Enough?" about DNS anycast at the DHS IMPACT PI meeting in San Diego, CA.

[Pre8] On 2015-02-10 John Heidemann gave the talk "New Opportunities for Research and Experiments in Internet Naming and Identification" at the CAIDA AIMS (Active Internet Measurement Systems) workshop in San Diego, CA. This talk introduced the concept of the research root in an public (although invited) workshop.

[Pre9] On 2015-02-26 John Heidemann and Terry Benzel briefed David Conrad about progress on B-Root Operations, and our goals for the research root. John Heidemann also gave the talk "Anycast Latency: How Many Sites are Enough?" about DNS anycast.

New in 2016q2:

[Pre10] On 2016-04-03, Wes Hardaker gave a talk on DNS RRL (Receiver Rate Limit) at the Root Operations meeting in Buenos Aires.

[Pre11]On 2016-03-11, John Heidemann gave a talk about the state of Internet Measurement Research at USC/ISI to Fred Baker and Murali Venkateshaiah from Cisco, exploring the possibility of industry support for our work.

New in 2016q3:

[Pre12] On 2016-07-17, Ricardo Schmidt gave the talk "Anycast Latency: How many sites are enough?" to the Root-Operators group in Berlin, Germany.

[Pre13]On 2016-07-17, John Heidemann gave the talk "Anycast vs. DDoS: Evaluating Nov. 30" to the Root-Operators group in Berlin, Germany.

[Pre14]On 2016-07-18, Givoane C. M. Moura gave the talk "Anycast vs. DDoS: Evaluating the Nov. 2015 Root DNS Event" to the IRTF MAPRG working group in Berlin, Germany.

[Pre15] On 2016-08-10, John Heidemann gave Matt Larson (CTO of ICANN) an update on B-Root status and recent work, including the analysis of the Nov. 30 event.

New in 2016q4:

[Pre16] On 2016-09-13, John Heidemann gave the talk "Anycast vs. DDoS: Evaluating Nov. 30" to the USC Networked Systems Laboratory. (Internal USC presentation.)

[Pre17] On 2016-10-11, John Heidemann gave the talk "Anycast Latency: How many sites are enough?" at an ICANN Lunch Seminar in Playa del Rey, CA, USA.

[Pre18]On 2016-10-16, John Heidemann gave the talk "Anycast Latency: How many sites are enough?" at the Fall DNS-OACR workshop in Dallas, Texas, USA.

[Pre19] On 2016-10-16, John Heidemann gave the talk "Anycast vs. DDoS: Evaluating Nov. 30" at the Fall DNS-OACR workshop in Dallas, Texas, USA.

[Pre20]On 2016-11-07, Wes Hardaker gave the talk "Critical Infrastructure DNS Research Testbed" at the ICANN DNSSEC Workshop in Hyderabad, India.

[Pre21]On 2016-11-15, Giovane Moura gave the talk "Anycast vs. DDoS: Evaluating Nov. 30" at the ACM Internet Measurements Conference in Santa Monica, CA, USA.

[Pre22]On 2016-11-17, John Heidemann gave the talk "Towards a Testbed for Naming and Internet Protocol Experimentation" at DINR-2016 in Marina del Rey, CA, USA.

[Pre23]On 2016-12-07, John Heidemann gave the talk “Distributed Denial-of-Service: What Datasets Can Help?” at Usenix ACSAC in Universal City, CA, USA, as part of the IMACT session chaired by Erin Kenneally.

6 Conclusions and Recommendations

6.1 Conclusions

We have completed all project deliverables, providing parallel analysis and playback of DNS data. In addition, we provided several results beyond original deliverables, including: analysis of anycast latency (as a technical paper), analysis of anycast under stress (as a technical paper), descriptions of a DNS testbed (as presentations). Finally, we have also released several software components as open source, including dnsanon and dnsanon_rssac.

6.2 Recommendations

Although we completed all objectives of the original proposal, there remains additional research to do. Our approach to DNS trace replay can be enhanced to support broader ranges of query types and query mutation. Our analysis of anycast can be enhanced to recommend corrective action in times of DNS stress. Our approaches to DNS anonymization are effective today, but further study of anonymization at different levels of the DNS hierarchy is warranted.

6. List of Symbols, Abbreviations and Acronyms

ACM	Association for Computing Machinery
ACSAC	Annual Computer Security Applications Conference
AFRL	Air Force Research Laboratory
AIMS	Active Internet Measurement Systems
AirOS	AirMAX Operating System
ANT	the Analysis of Network Traffic group (our research project, https://ant.isi.edu)
ARIN	American Registry for Internet Numbers, https://www.arin.net
AS	Autonomous System
BIND	Berkeley Internet Name Daemon (DNS server software)
CAIDA	Center for Applied Internet Data Analysis
CIO	Chief Information Officer
CTO	Chief Technology Officer
Co-PI	Co-Principal Investigator
DETER	The cyber Defense Technology Experimental Research Laboratory, a DHS-sponsored testbed, see https://deter-project.org
DHS	Department of Homeland Security
DINR	The DNS and Internet Naming Research workshop hosted at ISI.
DITL	Day-in-the-Life of the Internet, a roughly annual DNS data collection activity run by https://dns-oarc.net
DNS	Domain Name System
DNSSEC	Domain Name System Security
DNS-OARC	DNS Operations, Analysis and Research Center
DDoS	Distributed DoS
DoS	Denial of Service
DSC	DNS Stats Collector; software provide from https://dns-oarc.net
Dx.x	Deliverable x.x
ECMP	Equal Cost Multi-path
FIU	Florida International University
Gb/s	Gigabits per second
ICANN	Internet Corporation for Assigned Names and Numbers
IMC	the ACM Internet Measurement Conference
IMPACT	Information Marketplace for Policy and Analysis of Cyber-risk & Threats
IP	Internet Protocol
IPv4	IP version 4
IPv6	IP version 6
LACREND-RR	Los Angeles/Colorado Research Exchange for Network Data-Relay Root
LANDER	The packet collection software we developed; also the prior project funding the work: Los Angeles Network Data Exchange and

	Repository
LAX	Los Angeles International Airport
LDplayer	DNS trace replay software we are developing; see https://ant.isi.edu/software/ldplayer/
Mb/s	Megabits per second
OCSP	Online Certificate Status Protocol
PAM	Passive and Active Measurements Conference
PI	Prototype Improvement and Principle Investigator
PREDICT	DHS Protected REpository for the Defense of Infrastructure against Cyber Threats (PREDICT) project.
PreX	Presentation X
PubX	Publication X
RRL	Receiver Rate Limiting
RSSAC	Root Server System Advisory Committee
STx.x	Subtask
TCP	Transmission Control Protocol
TI	Texas Instruments
TTx	Technology Transfer
USC/ISI	University of Southern California/Information Sciences Institute